

Report of the Interagency
Optical Network Testbeds Workshop 2
NASA Ames Research Center

September 12-14, 2005



Jointly sponsored by the
Department of Energy Office of Science
National Aeronautics and Space Administration

Joint Engineering Team
of the Networking and Information Technology R&D Program's
Large Scale Networking Coordination Group

National Coordination Office for Networking and Information Technology Research and Development

This document was prepared and published by the National Coordination Office for Networking and Information Technology Research and Development (NCO/NITRD). The NCO/NITRD supports overall planning, budget, and assessment activities for the Federal multi-agency NITRD Program under the auspices of the NITRD Subcommittee of the National Science and Technology Council.

Copyright Information

This is a work of the U.S. Government and is in the public domain. It may be freely distributed, copied, and translated; acknowledgement of publication by the National Coordination Office for Networking and Information Technology Research and Development is appreciated. Any translation should include a disclaimer that the accuracy of the translation is the responsibility of the translator and not the NCO/NITRD. It is requested that a copy of any translation be sent to the NCO/NITRD.

To Request Additional Copies

To request additional copies of the *Report of the Interagency Optical Network Testbeds Workshop 2* or other NITRD Program publications, please contact: NCO/NITRD, Suite II-405, 4201 Wilson Boulevard, Arlington, Virginia 22230; (703) 292-4873; fax: (703) 292-9097; e-mail: nco@nitrd.gov. Electronic versions of NITRD documents are also available on the NCO Web site: <http://www.nitrd.gov>.

Disclaimer

Companies and products identified in this report were cited in discussions that took place at the Optical Network Testbeds Workshop 2 and/or reflect partnerships that exist in the optical networking community. Such identification does not imply recommendation or endorsement by the Federal government, nor does it imply that the companies or products referred to are the best available for an intended purpose.

ONT2 Workshop Report

INTRODUCTION

The Federal networking research and development (R&D) community, whose visionary technical breakthroughs led to the Internet and the information technology (IT) economy, continues to develop new evolutionary – and revolutionary – technologies that drive broad networking advances as well as address pressing Federal agency mission needs. Recent Federal R&D efforts have increased network bandwidth, reliability, security, and differentiated services; enabled dynamic, hybrid, and mobile networking; and created new paradigms for distributed capabilities such as grid networking and sensornets. These efforts, often in partnership with industry, have been adopted rapidly in the marketplace, providing foundations for further expansion of the IT economy.

Today, Federal networking R&D is increasingly focused on the revolutionary promise of optical technologies. These new photon-based technologies will enable next-generation networks that are orders of magnitude faster, more secure, more flexible, and higher-bandwidth than the current Internet. Such capabilities are needed to support Federal agencies' science and engineering research, national defense, and national security missions, which increasingly require rapid, safe, and reliable transmission of petabytes of real-time data, support for extremely data-intensive distributed applications and infrastructure, and other advanced services not possible over current networks.

Federal R&D in advanced optical networking technologies is also demonstrating that these powerful new capabilities will enable the Nation's IT infrastructure to continue to evolve without disruption. The IT infrastructure's current generation of interconnected computers, storage systems, software, data resources, mobile devices, sensor systems, and other digital equipment – encompassing not only the Internet but the vast range of public- and private-sector enterprise networks that connect to it – now faces constantly expanding technical challenges of interoperability, reliability, security, and scalability across diverse legacy technologies. The optical technologies now being prototyped leapfrog these limitations, making possible wholly new paradigms of networking architecture that provide enormous flexibility for differentiated services (e.g., highly secure, dynamic mobile networking, bulk data transfers, sensornets) while accommodating existing IT infrastructure and services.

To accelerate the work of the optical networking R&D community, Federal R&D agencies in cooperation with university and industry partners are building optical network testbeds (ONTs) that develop and demonstrate new networking concepts, architectures, and technologies on real networks supporting real users with real applications. In addition, the optical networking researchers are participating in a series of workshops designed to expand their collaborative activities, shape the overall optical networking research agenda, and develop a shared vision for the optical network future across the national and international networking R&D communities.

At the first ONT workshop, Federal network program managers, optical network researchers and developers, and companies with an interest in networking research developed plans to coordinate research among the several testbeds; ensure that technology developed in each testbed would be shared in a timely fashion among all the testbeds to foster standard technologies and middleware; and support ongoing testbed coordination. The ONT2 workshop, summarized in this report, discussed progress on the ONTs and identified opportunities for further coordination and collaboration. The next workshop, ONT3, to be held in Tokyo, Japan on September 7-8, 2006, will expand the ONT coordination and collaboration to the international arena.

ONT2 Workshop Report

TABLE OF CONTENTS

Introduction	iii
Table of Contents	v
1.0 Executive Summary	1
1.1 Workshop Overview	1
1.2 Vision: Enabling the Optical Networking Future	2
1.3 Steps to Accelerate R&D Progress	2
1.4 Steps to Transfer R&D Results to Operational Networks	3
2.0 Workshop Report	5
2.1 Current Focus on Optical Networking R&D by Networking Organizations	5
2.2 Identification of ONT Research Areas	6
2.3 Interoperability	7
2.3.1 Current Status of Interoperability Research	7
2.3.2 Recommended Actions for Interoperability	8
2.4 Common Internet Service Provisioning (CISP)	9
2.4.1 CISP Research Areas	10
2.4.2 Recommended Actions to Promote CISP	10
2.5 Applications Support	11
2.5.1 Current ONT Support for Applications	11
2.5.2 User Requirements	12
2.5.3 Recommended Actions on Applications Support	12
2.6 Open Exchange Points (OEPs)	13
2.6.1 Recommended Actions on OEPs	13
2.7 Core Needs of the Research Community	14
2.7.1 Scalable Testbed	14
2.7.2 Coordination on Emerging Architectures	14
2.7.3 Persistent Funding	14
2.7.4 Recommended Actions on Research Community Needs	14
2.8 Operations and Management	15
2.9 Cost Models and Building Funding Support	16
2.9.1 Recommended Actions on Cost Models and Funding Support	17
2.10 MorphNet Architecture to Support Transition From R&D to Production	17
Appendix A: Workshop Administration and Attendees	18
Appendix B: About LSN and the NITRD Program	21
Abbreviations/Acronyms	23

1.0 EXECUTIVE SUMMARY

1.1 Workshop Overview

The Optical Network Testbeds Workshop 2 (ONT2), held on September 12-14, 2005, was co-sponsored by the Department of Energy Office of Science (DOE/SC) and the National Aeronautics and Space Administration (NASA), in cooperation with the Joint Engineering Team (JET) of the Federal Networking and Information Technology Research and Development (NITRD) Program's Large Scale Networking (LSN) Coordinating Group (see "About LSN and the NITRD Program" on page 21).

The ONT2 workshop was a follow-on to an August 2004 Workshop on Optical Network Testbeds (ONT1). ONT1 recommended actions by the Federal agencies to assure timely development and implementation of optical networking technologies and infrastructure (see <http://www.nren.nasa.gov/workshop7>).

Hosted by the NASA Ames Research Center in Mountain View, California, the ONT2 workshop brought together representatives of the U.S. advanced research and education (R&E) networks, regional optical networks (RONs), service providers, international networking organizations, and senior engineering and R&D managers from Federal agencies and national research laboratories. Its purpose was to:

- Develop a common vision of the optical network technologies, services, infrastructure, and organizations needed to enable widespread use of optical networks
- Recommend activities for transitioning the optical networking research community and its current infrastructure to leading-edge optical networks over the next three to five years
- Present information enabling commercial network infrastructure providers to plan for and use leading-edge optical network services in that time frame

Background information considered by the ONT2 participants included a 5-to-15-year network research perspective presented in a report on two recent National Science Foundation (NSF) workshops that examined long-term research goals and directions. This report, *Mapping a Future for Optical Networking and Communications*, by Dan Blumenthal, John Bowers and Craig Partridge, is available at: <http://www.geni.net/nsf-opt-200507.pdf>.

Participants discussed the workshop's goals in the following panel sessions: innovative prototype networks; basic optical network research testbeds; production R&D networks, including networks supporting Federal network research and applications requiring advanced networking, and selected regional optical networks (RONs); international R&D networks and commercial network technology and service providers; and long-term views based on the NSF report. The participants concluded that the ONT2 discussions effectively surfaced multiple key ideas, research needs, and organizational issues that provided a starting point for future progress reviews and roadmapping activities.

Summary reports from all the panels are available at <http://www.nren.nasa.gov/workshop8>. The industry perspectives presented at the workshop are available at: <http://www.nren.nasa.gov/workshop8/index.html>.

ONT2 Workshop participants also developed plans for the next ONT workshop. The Japanese National Institute of Information and Communications Technology (NICT) agreed to host

ONT2 Workshop Report

Optical Network Testbeds Workshop 3 (ONT3) in Tokyo on September 7-8, 2006, in coordination with the 2006 meeting of the Global Lambda Integrated Facility (GLIF), an international virtual organization that promotes lambda networking (e.g., wavelength-based services) to support data-intensive scientific research and middleware development. ONT3 will enable LSN agency senior testbed and network managers to evaluate their progress in carrying out the ONT2 recommended actions over the upcoming year and to coordinate their ongoing programs with counterparts from government-funded research testbeds and R&E networking communities worldwide.

1.2 Vision: Enabling the Optical Networking Future

The optical networking future envisioned by the ONT2 panels will offer new networking services and technologies as core components of next-generation networks. For example, Layer 1/Layer 2 networking adds a new “underlay” of services to existing internetworking (compared with Multi-Protocol Label Switching [MPLS], which adds an overlay) that will enable innovative and flexible network architectures. Direct light paths will allow some incoming high-performance connections to bypass the bottlenecks of Internet routers and firewalls, when required, and to connect directly to campus-based high-performance data sources, supercomputing systems, and visualization users. New protocols will enable higher speeds over longer distances, using either a single light path, multiple dedicated light paths, or an optical private network (OPN). Such networks and protocols will increase user control by providing users the ability to enable applications to request, utilize, and tie together high-performance services deep within the network.

1.3 Steps to accelerate R&D Progress

To accelerate research progress toward achieving this vision, ONT2 participants identified the following areas in which the various optical networking R&D communities need to engage and collaborate:

Testing real applications on an entirely new kind of network

The advent of new types of Layer 1/Layer 2 optical networking services over the next few years will enable networks to address applications performance problems over long distances and covering multiple networking domains. The LSN and partner communities – particularly the JET – should coordinate the research prototyping testbed networks to address the performance problems of applications scientists funded by the Federal agencies. These services will also enable new applications (e.g., those requiring precise service guarantees) that are not possible using traditional networks.

Acting to obtain to dark fiber

The R&D to enable this vision in the next few years requires access to many miles of dark (not in use) fiber – whether agency-owned or shared with other agency, national, or regional networks – and/or access to light paths within lit (in use) fiber. The window of opportunity to buy dark fiber to support long-term networking R&D at today’s reasonable prices is closing rapidly as existing fiber capacity is increasingly utilized and prices increase. The U.S. National LambdaRail (NLR) and the regional optical networks have already purchased low-cost dark fiber and demonstrated the improved cost-effectiveness of this approach. Federal agencies and partner organizations should take advantage of these low costs while they still exist.

ONT2 Workshop Report

Making changes at the campus level

With the advent of optical networks, Federal laboratory and university network managers will need to manage the integration of optical networking infrastructure into the existing campus network infrastructure. Currently, multiple optical networks coming to campus sites cannot be managed at the campus level. Campus sites may need to build a fiber infrastructure to a RON to enable optical network access at the campus level.

Participating in large-scale initiatives

The optical networking R&D communities should actively support and contribute to large-scale efforts aimed at advancing optical networking technologies. Two such efforts identified by the ONT2 participants are:

GLIF: The Global Lambda Integrated Facility (GLIF) (www.glif.is) brings together senior networking engineers to design and develop an integrated international lambda infrastructure by identifying equipment, connection requirements, and needed engineering functions and services. These capabilities are currently offered through a global distributed infrastructure based on advanced optical networking. The GLIF participants are national research and education networks (NRENs), consortia, and institutions working with lambdas. Administrative support is provided by TERENA with financial support from sponsoring organizations.

GENI: The Global Environment for Networking Investigations (GENI) (<http://www.nsf.gov/cise/geni>) sponsored by NSF is a long-term initiative to explore new networking architectures for advancing science and stimulating innovation and economic growth. The GENI initiative plans a research program and a global experimental facility designed to explore new network architectures at scale. It will explore creating new core functionality, deploying and validating new network architectures, building higher-level service abstractions, and developing new services and applications. The GENI facility will create a large-scale experimental infrastructure to support the initiative's research goals. GENI is envisioned as a broad community effort engaging other agencies, other countries, and corporate entities.

1.4 Steps to Transfer R&D Results to Operational Networks

ONT2 participants also identified actions needed or already under way to develop and promote the transfer of optical networking technology, services, and infrastructure from the R&D arena to the operational networking community. The action areas include:

Working toward multi-domain interoperability

- Providing end-to-end optical networking across even two or three domains is currently a challenge both at the control plane layer and at higher layers. However, network researchers are making progress and the work should be continued and expanded. Development of a multi-domain control plane will take place over several years through bilateral, inter-testbed, and inter-regional efforts. Prototype methods should be coordinated globally through the GLIF and the Global Grid Forum (GGF, <http://www.gridforum.org>).
- Senior network engineers should work through events such as iGRID (annual international grid workshop) and within the GLIF to agree on common sets of functionalities for Open Exchange Points (OEPs) to support interoperable bandwidth requests, inter-domain management, and other service requests through facilities called Global Open Lambda Exchanges (GOLEs).

ONT2 Workshop Report

- Federally supported optical networks should participate with DOE's Energy Sciences Network (ESnet) On-demand Secure Circuits and Advanced Reservation System (OSCARS) project and the Internet2 Bandwidth Reservation for User Work (BRUW) project, in coordination with GÉANT2 (the pan-European research and education network) and CANARIE (Canada's advanced Internet development organization), to coordinate on intra- and inter-domain dynamic provisioning and routing.

Agreeing on common service definitions

- To support key user services, testbeds such as NSF's Dynamic Resource Allocation via GMPLS Optical Networks project (DRAGON) should develop engineering specifications (including cost/price as appropriate) followed by encodings in formal languages such as Backus-Naur Form (BNF) or eXtensible Markup Language (XML).
- The optical networking research testbed community should collaborate with the GLIF community to define and adopt basic light path service definitions and a standard architectural context for those services, such as the Web Services Resource Framework (WSRF).
- The R&D community should develop several operational optical OEPs with services based on agreed common service definitions and should review progress at the ONT3 workshop.

Developing a large-scale multi-year research testbed

- The research community should involve research participants from related knowledge domains in a large-scale optical research testbed.
- The research community should define funding models for multi-year support of key research objectives.

Promoting distributed applications and use of Layer 1/Layer 2 services

- Testbeds should make additional light paths available to selected applications and work with the researchers on how to use lambdas to support their applications.
- DoD's Global Information Grid Evaluation Facilities (GIG-EF) and NSF's DRAGON are working together on interoperable networking for the Massachusetts Institute of Technology's Haystack Observatory astronomy facilities and will share their approaches and experiences with others.
- The Hybrid Optical Packet Internet (HOPI) project of Internet2 (the advanced networking consortium of academic institutions) and DOE/SC's UltraScience Network (USN) have agreed to collaborate on an application yet to be identified.
- GIG-EF and USN researchers will consider collaborating on distributed applications using the open-source LUSTRE (acronym from "Linux" and "cluster") distributed file system.

2.0 WORKSHOP REPORT

The Optical Network Testbeds Workshop 2 (ONT2) brought together members of the U.S. research and education (R&E) networks of the NITRD Program's Large Scale Networking Joint Engineering Team (JET), partner advanced networks, regional optical networks (RONs), service providers, international networking organizations, and senior engineering and R&D managers from Federal agencies and research laboratories. Its purpose was to develop a common vision of the optical network technologies, services, infrastructure and organizations needed to enable widespread use of optical networks and to use this vision to develop recommended activities for transitioning the optical networking research community and its networking infrastructure to utilizing leading-edge optical networks over the next three to five years.

In six breakout panels and a concluding plenary session, workshop presenters identified their current optical network testbed research and development (R&D) programs, key issues that need to be addressed, technology-specific roadmaps for development and implementation, and recommendations for Federal actions to promote the use of optical networking in operational networks. Participants concluded that the ONT2 discussions effectively surfaced multiple key ideas, research needs, and organizational issues that provide a starting point for future progress reviews and roadmapping activities.

2.1 Current Focus on Optical Networking R&D by Networking Organizations

The majority of optical networking R&D activities are taking place on experimental testbeds, designed specifically for research. Additionally, some networking organizations are exploring optical networking technologies for production services. Current research interests of networking organizations with optical facilities include:

CA*net4, Canada's optical research and education network, is providing users with more control to manage, deploy, and use optical networks through implementation of service-oriented-architecture (SOA) techniques and high level work flow processes.

HOPI, Internet2's Hybrid Optical Packet Internet, is developing operations that serve all layers of the protocol stack, working on inter-domain interoperability across optical networks, and providing support for new applications.

USN, DOE/SC's UltraScience Network, is developing and inter-coordinating gigabit Ethernet/MPLS interoperability, services, and security.

SURFnet, the Netherlands' research and education network, is focused on global end-to-end interoperability, access to dark fiber for the backbone and the customer, inter-domain interoperability, protocols, user support, and implementing exchange points including those using advanced optical technologies such as NetherLight.

GIG-EF, DoD's Global Information Grid Evaluation Facilities, is coordinating the building of the control plane, Session Initiated Protocol (SIP), Web services, pervasive native IPv6, and end-to-end QoS/Priority for high-end information streams.

ONT2 Workshop Report

StarLight, the NSF-funded international optical exchange point, has implemented prototype Layer 1/Layer 2 based services and made them available to the national and international research community and global science communities.

National LambdaRail (NLR), a research network consortium, has a national fiber infrastructure facility providing multiple wavelengths for the research community – one-half of its capacity is allocated for networking R&D projects. NLR managers describe it as “not a single network but a set of facilities, capabilities, and services to build both experimental and production networks at various layers, allowing members to acquire dedicated (project-specific) facilities or shared (community-specific) facilities as appropriate.”

2.2 Identification of ONT Research Areas

Workshop participants identified R&D needed to foster the operational use of optical networking. Interoperability among ONT domains is not only a current interest of the ONT testbeds but also a key area for future R&D, especially Common Internet Service Provisioning (CISP), which provides end-to-end services across domains and standards to foster transparency. There was significant overlap on issues identified by the workshop panels and the actions they proposed. The panels were:

- Panel A: Innovative prototype networks
- Panel B: Basic optical network research testbeds
- Panels C and D: Production R&D networks, including networks supporting Federal network research and applications requiring advanced networking and selected RONS
- Panel E: Long term, 5-to-15 year views based on two NSF workshop reports
- Panel F: International R&D networks and commercial network technology and service providers

Panelists identified the following topic areas as research priorities:

- **Interoperability** including:
 - Optical transport
 - Data plane
 - Control plane
 - Management plane
 - Routing plane
 - Service plane
 - Technology for providing fault isolation
- **Common Internet Service Provisioning** including:
 - Providing a common, interoperable service to respond to user requests
 - Common service definitions
- **Applications Support** including:
 - Services

ONT2 Workshop Report

- Service-oriented architecture
- Tools enabling users to manage and provision their own networks
- Capabilities that stimulate new applications development on these networks
- Middleware that provides applications with new addressable communications functionality
- **Open Exchange Points** including:
 - Well-defined lower layers
 - Framework and standards
 - Tools for provisioning
 - Dynamic provisioning
 - Identification and mitigation of risks
- **Core Research Interests** including:
 - Common service definitions
 - Experimental optical network infrastructure
 - New architectures for control and management
 - Persistent (multi-year) research funding
- **Operations and Management:** The “care and feeding” of users, applications, and the network while maximizing service offerings
- **Cost Models:** To ensure the network’s financial sustainability
- **MorphNet Architecture:** A concept developed at DOE’s Argonne National Laboratory proposing a multi-service architecture within a shared infrastructure

The panels’ findings and recommended actions in these areas are discussed in the following sections.

2.3 Interoperability

2.3.1 Current Status of Interoperability Research

Workshop participants provided briefings on the current status of interoperability across the ONT testbeds and roadmaps of technologies, services, and applications needed to promote use of optical networking (see: <http://www.nren.nasa.gov/workshop8/>). The current status of ONT testbed interoperability includes:

HOPI, DRAGON, and CHEETAH

HOPI, DRAGON, and CHEETAH (NSF-funded Circuit-switched High-speed End-to-End Transport ArchHitecture project) are working on interoperability and transparency. HOPI is also working with regional networks and international networks. It has implemented an OC-192 link to Europe for connection to London and Amsterdam that is currently being used for demonstrations and experiments. The HOPI beta control plane is in place and providing support for the DOE/SC- and NSF-sponsored UltraLight project (physics applications) at the California Institute of Technology. It has not been open to other users.

ONT2 Workshop Report

The HOPI Testbed Support Center has three components:

- The engineering component, provided by the Indiana University network operations center (NOC), is operational.
- The control plane component is in beta testing.
- The applications component supports existing applications and promotes creative thinking for developing new applications over various networking paradigms.

UltraScience Net

USN interoperability has three objectives:

- Provide a control plane user interface
- Interoperate with CHEETAH
- Support CHEETAH interoperability with DRAGON and HOPI

Global Information Grid Evaluation Facilities

The GIG-EF is collaborating on research with the Mid-Atlantic Exchange (MAX), ATDnet (using the same fibers), DRAGON, and BOSSnet to provide connectivity to the Haystack Observatory. These communities currently switch wavelengths among these networks (using optical switching implemented with GMPLS) and with the Naval Research Laboratory at the physical layer, data plane, and control plane. GIG-EF will work with BOSSnet and DRAGON to implement SIP among these networks.

SURFnet

SURFnet is collaborating within the GLIF to implement data plane (Layer 1) interoperability among Netherlands (NetherLight), Chicago (StarLight), Seattle (Pacific Wave), and New York (Manhattan Landing [MAN LAN]). GÉANT2 (the pan-European research and education network) also expects to be fully compatible with these networks at Layer 1 within the next year.

These organizations have an objective of reaching agreement on global approaches to control plane interoperability (interoperability within a domain is achieved by using a single vendor so the issue is interoperability between and among domains). It would be premature to establish global standards at present, but coordination with the GLIF Control Plane Working Group (WG) and the GGF is needed to develop a global approach to control plane interoperability, rather than a proliferation of national or regional approaches.

At higher network layers, ONT2 participants said, control plane interoperability should be achieved based on Ethernet framing. At Layer 2, lambdas should be connected to central points, e.g., to the MAN LAN Ethernet switch, where networks would establish bilateral agreements for peering. At higher layers, the SIP might be used, e.g., to request Layer 2 and Layer 3 services for video applications.

2.3.2 Recommended Actions for Interoperability

- *ONT providers should work with existing GLIF activities to implement global interoperability. Interoperability at the data plane is a mature technology.*

ONT2 Workshop Report

- *Multi-domain control plane interoperability is a multi-year problem that will most likely be solved in an evolutionary way by bilateral inter-testbed and inter-regional efforts. All testbeds should collaborate globally in major forums such as GLIF and GGF.*
- *Provide common service definitions: See the discussion in 2.4 below.*
- *Testbeds should make light paths available to selected applications and should work with the applications developers to effectively use lambdas for those applications.*

2.4 Common Internet Service Provisioning (CISP)

Service provisioning allocates or shares network resources, including instantiating network services, within or across autonomous domains or vendors. In the past, intra- and inter-domain provisioning was not implemented across network boundaries for technical and economic reasons, despite the existence of enabling technologies such as Asynchronous Transfer Mode (ATM). Today, modifications and extensions to existing boundary protocols (e.g., Border Gateway Protocol [BGP]) allow for delineation of network boundaries and the possibility of secure inter-domain provisioning. Current inter-domain signaling protocols such as MPLS may also provide a basis for intra- and inter-domain provisioning.

For successful (manual or dynamic) inter-domain provisioning, all networks located along the end-to-end path must share a clearly defined common service offering. This applies to all levels of inter-domain hierarchical connections: RON-backbone, RON-RON, and RON-campus. Workshop participants identified the GLIF as the proper forum for the creation and maintenance of the service definitions. They indicated that the Quilt (a U.S. research and education networking coalition also concerned with inter-domain service definitions) would be likely to follow the lead of the GLIF.

The ESnet OSCARS project and the Internet2 BRUW project, in cooperation with GÉANT2 and CANARIE, are currently developing intra- and inter-domain dynamic provisioning and routing. Where appropriate, software developed by these projects will be made available to other networks.

ONT2 participants identified the need to securely allocate bandwidth and/or light paths within or across autonomous domains while efficiently sharing scarce network resources. The domains need to interoperate and enable these services across domains with repeatable, known, consistent characteristics.

Global interoperability requires cross-domain provisioning of Layer 1 and Layer 2 communications services, such as optical light paths and Ethernet framed service. This requires specific definitions of these services so that providers can translate them to their networks to set up the specific services requested by users. Currently, there are many parameters left undefined when a user requests a service. For example, “100 Mbps Ethernet service from point A to point B” could be interpreted in several ways, including:

- Half duplex, or full duplex

ONT2 Workshop Report

- Tagged or untagged virtual local area networks (VLANs)
- Byte-frame size of 1,500, 4,470, or 9,216 bytes
- Full or shared bandwidth

Common service definitions and formalized descriptions are needed so that service options can be presented and understood by all parties, from the service provider to the end consumer. R&E networking architects and engineers should formalize these definitions. Specifications need to be developed that customers can measure, and engineers can test to, the service levels agreed upon during negotiation and contract phases. A white paper discussing the justification and associated issues for “Common Service Definitions” (Sobieski/Lehman) can be found at: <http://dragon.maxgigapop.net>.

At the GLIF meeting immediately following the ONT2 workshop, the GLIF Control Plane WG continued its work on automating end-to-end connections over GLIF resources, with the ultimate goal of using control plane signaling and restoration. Representatives of the DRAGON/MAX project gave a presentation to the WG entitled “Common Service Definitions.” CANARIE representatives presented a briefing on the usefulness of SOAs and their potential role in automation. Abilene representatives gave a presentation on testing and verification of services. The WG identified application parameters needed to support common service definitions including jitter and latency. WG members also noted that, for many services, users will need mechanisms to verify end-to-end service delivery. A GLIF task force led by DRAGON researchers will develop an initial set of common service definitions.

2.4.1 CISP Research Areas

ONT2 workshop participants identified the following topic areas as priorities for CISP R&D:

- Intra- and inter-domain provisioning
- Resource sharing
- Dynamic and static networking
- Lambda allocation
- Prioritization and preemption of traffic
- MPLS tunnels
- User-controlled light paths (UCLPs)
- Security and authentication, authorization, and accounting (AAA)
- Testing and inter-domain fault detection

2.4.2 Recommended Actions to Promote CISP

- *Testbeds such as DRAGON should identify key user services, develop engineering specifications for them (including cost/price as appropriate), and encode them in appropriate formal languages, e.g., BNF or XML.*
- *ONT researchers should use open venues like iGRID, SuperComputing (SC), GGF, and GLIF meetings to discuss key user services and the developments taking place within the Internet Engineering Task Force (IETF), International Telecommunications Union (ITU), Optical Internetworking Forum (OIF), Quilt, GGF, and GLIF.*

ONT2 Workshop Report

- *The community should define and adopt an initial set of basic light path service definitions working within the GLIF framework.*
- *The community should implement several OEPs based on agreed common service definitions and review progress in this area at ONT3 in September 2006.*
- *Research networks, the ESnet OSCARS project, the Internet2 BRUW project, G ANT2, and CANARIE should cooperate on developing intra- and inter-domain dynamic provisioning and routing.*
- *The JET should identify actions to promote CISP and work to implement them.*

2.5 Applications Support

2.5.1 Current ONT Support for Applications

Workshop participants from the following ONTs described their current applications support programs:

HOPI

The HOPI Testbed Support Center supports applications users through its NOC at Indiana University and its applications component, which assists existing applications and provides consultation in developing new applications over various networking paradigms. HOPI is building sets of equivalence classes of networking services for support of the new applications and plans to have these services working cooperatively to support new applications within one year.

GIG-EF

GIG-EF strategy for encouraging and supporting new applications includes:

- Building applications that can decide whether to use standard grid technology versus using InfiniBand
- Moving InfiniBand outside the computer room to the wide area network (WAN)
- Providing distributed file systems on the desktop

UltraScience Net

USN supports four key applications:

- Remote control of a high-speed optical confocal microscope, which generates high-resolution imagery with requirements for low latency and low jitter. This application streams image data for analysis, which is rendered into video products and then streamed to users.
- Lambda Station – A high-energy physics application developed as part of the Large Hadron Collider (LHC) effort, using intelligent data steering to reserve USN bandwidth and make optimal use of it.
- A dedicated data and control channel-based visualization application being developed for the DOE/SC-supported Terascale Supernova Initiative. This application uses the USN

ONT2 Workshop Report

ability to create parallel, independent dedicated channels at specified bandwidths between the same two endpoints.

- LUSTRE – An open source, geographically distributed file system for supercomputers being developed at several centers. USN is serving as a testbed for extending LUSTRE to the wide area, using the USN Layer 2 ability to operate without router-generated session delays.

2.5.2 User Requirements

User requirements for applications support are increasing because applications hardware is getting cheaper while running at higher data rates. For example:

- Tile displays now available for \$20 K can consume 5 Gbps or more bandwidth, which requires a dedicated lambda and high-quality service guarantees to run transparently.
- Single low-cost end devices can now generate or consume up to 8 Gbps.
- High-end COTS servers available for \$10 K can burst at 30 Gbps.

Such improved hardware technologies enable more data-intensive applications, increasing the burden on networks to support the large data transfers that these applications require. Some approaches identified by workshop participants to accommodate the increased requirements include:

Open source middleware

Open source middleware is needed to expedite the sharing of developing ONT capabilities, e.g., resource brokers and standards.

Service-oriented architectures

SOAs support the development of clean application program interfaces (APIs) at key levels for ONT users and providers. Using SOAs, DRAGON plans to develop an application-specific topology enabling an application request to quickly establish a target topology at multiple layers. Collaboration with partners to discuss research approaches, exchange notes, and compare efforts would promote multi-layer, multi-domain interoperability using SOAs.

Increased collaboration

Interoperability among Federal agency facilities, users, and university researchers would benefit from increased collaboration among Federal research networks, e.g., DRAGON, CHEETAH, USN, HOPI, and GIG-EF communities. Renewed collaboration is needed among DARPA, NSF, NASA, and DoD networks, such as through holding joint principal investigator meetings.

2.5.3 Recommended Actions on Applications Support

- *It is premature to implement global multi-domain interoperability, but progress is being made by the ONTs on inter-domain interoperability among a limited number of domains. This work should be continued and collaborative efforts should be expanded.*

ONT2 Workshop Report

- *The networking R&D communities should use opportunities such as the annual iGrid and SC conferences to move forward incrementally on working together through discussion, information exchange, and adoption of improving approaches and practices.*
- *GIG-EF and DRAGON, currently working together on Haystack Observatory facilities interoperability, will share their approaches and experiences with the community.*
- *HOPI and USN should identify an application on which they will collaborate.*
- *GIG-EF and USN should consider collaborating on supporting distributed applications using the LUSTRE distributed file system.*

2.6 Open Exchange Points (OEPs)

GLIF initiatives have fostered the implementation of OEPs at Layer 1. Following the GLIF architecture, they are labeled Global Open Lambda Exchanges (GOLEs). The current GOLEs include:

- | | |
|---|---------------------------------------|
| ▪ CANARIE – StarLight (Chicago) | ▪ NetherLight (Amsterdam) |
| ▪ CANARIE – Pacific Northwest GigaPoP (Seattle) | ▪ NorthernLight (Stockholm) |
| ▪ CENIC LA PoP (Los Angeles) | ▪ Pacific Northwest GigaPoP (Seattle) |
| ▪ CERN (Geneva) | ▪ StarLight (Chicago) |
| ▪ KRLight (Seoul) | ▪ T-Lex (Tokyo) |
| ▪ MAN LAN (New York) | ▪ UKLight (London) |

In addition :

- NGIX-East in the Washington, D.C., area, which currently switches inter-testbed traffic at Layer 1, soon may be advertised as a Layer 1 OEP.
- SOX in Atlanta and AMPATH in Miami are working on OEPs.

With the right framing, such as Ethernet, OEPs automatically support bilateral peering agreements at Layer 2. OEPs typically provide IP/BGP support for Layer 3 as well. There is interest in the U.S. in developing distributed OEPs (e.g., Pacific Wave has Ethernet/Nortel optical cross connects [HDXcs] in Los Angeles and Seattle interconnected with a dedicated E-wave).

2.6.1 Recommended Actions on OEPs

- *Senior network engineers should work through events such as iGRID and within GLIF on common sets of functionalities for OEPs.*
- *Sponsoring organizations should “think globally, act locally,” by providing OEPs for the community at key points of global interconnection.*
- *Federal agencies should promote research on:*

ONT2 Workshop Report

- *Middleware*
- *AAA to support peering*
- *Call blocking*

2.7 Core Needs of the Research Community

ONT2 workshop participants identified the following core research community needs for fostering development and use of optical networking.

2.7.1 Scalable Testbed

The research community requires an experimental optical network testbed that scales, over the long term, at both the infrastructure and community levels. This testbed would help solve challenging problems posed by promising but difficult-to-develop technologies that could be highly disruptive. The testbed would be used to investigate methods, architectures, systems, subsystems, and advanced optical devices under development.

The testbed should involve larger communities of interest such as leading-edge science communities that traditionally drive development of advanced technology. To engage a wider range of research communities, partnerships should be developed with the advanced technology developers of the TeraGrid and data grids as well as with Federal agencies, research centers, and international testbeds. This goal requires implementing techniques for bringing these new technologies, such as light path services, to end users. Currently, because of their security firewalls and lack of tools to support optical networking, campus networks typically constitute a barrier to advanced optical services.

2.7.2 Coordination on Emerging Architectures

The research community should use GLIF to coordinate emerging new architectures for distributed management and control planes. Basic research issues include developing common definitions, inter-domain and application policy systems, instantiation methods at multiple levels (e.g., network, operations, business), interoperability methods, and measurement techniques (e.g., for security, data integrity, performance, and validation).

2.7.3 Persistent Funding

Stable funding over several years is needed to achieve basic research goals and to regularly refresh optical testbed technology. ONT2 workshop participants also recommended that agencies allocate research funding for interdisciplinary projects, which may require policy changes at some agencies that primarily support existing discipline areas. In addition, participants recommended allocation of funding for innovative research projects among testbeds. (See additional discussion of funding in section 2.9, Cost Models.)

2.7.4. Recommended Actions on Research Community Needs

- *The optical networking testbed research community should identify opportunities for collaboration on service definitions with the GLIF community.*

ONT2 Workshop Report

- *Optical networking researchers should define the objectives, processes, operations, and support models for a large-scale optical research testbed involving research participants from other investigative knowledge domains.*
- *The optical networking testbed research community should explore opportunities for collaboration, especially with the GLIF community, on the design and development of next-generation management and control planes. Innovative cooperative research projects among international testbeds represent an important opportunity.*
- *The research community should develop funding models to support the key objectives over a multi-year timeframe.*

2.8 Operations and Management

The area of operations and management supports the network fabric and users to maintain service offerings. Optical networking operations and management components include monitoring network activity, understanding and managing traffic patterns, quick detection and resolution of error conditions, and testing. As service offerings become more varied and complex (and use multiple layers of the protocol stack), fault isolation is expected to become a more complex challenge. In addition to locating the “break in the pipe,” fault isolation will need to identify mismatched descriptions of services resulting in leaks and obstructions.

Monitoring network performance to quickly and efficiently isolate network faults requires tools that allow visualization of provisioned paths or virtual services from a high level, with the ability to “drill down” to obtain specific performance statistics or to analyze problems. Statistics are usually available from the command-line interface of the network devices, but a graphical tool could provide a better user experience. Vendors should provide better public APIs so that network operators can develop their own monitoring tools. The NSF End-to-end Diagnostic DiscoverY (EDDY) project monitoring tool collects data on application, network, system, environmental, and security events from multiple subsystems, normalizes them into a common event record (CER), and disseminates them into a highly distributed diagnostic backplane where diagnosticians can perform both general and highly focused forensic analysis on a wide variety of faults and anomalies.

Currently, research networks have differing levels of acceptance testing and methods for monitoring network health and performance. However, there are cooperative interagency network performance projects. Workshop participants suggested that sharing network test equipment, particularly 10 Gbps equipment, could reduce the overall costs of network testing. Today, many of the performance testing tools used by the research community are commonly used on routed networks. For example, specific Layer 3 projects that test application-level performance over networks include Abilene’s use of NetFlow data for flow performance analysis. In addition, the Abilene network is performing H.323 (a proposed set of audio-visual communication protocols recommended by the ITU) beacon testing by injecting H.323 streams into the network and analyzing the results. Although these tools work for routed networks, many Layer 3 techniques are not useful for Layer 1/Layer 2 analysis. However, existing forums could be used for information sharing on new Layer 1/Layer 2 methods. For example, new

ONT2 Workshop Report

technologies like SFlow (implemented by some vendors) have the capability of monitoring Layer 2 flows through a network device and can sample IPv6 as well as IPv4 traffic. Initiatives such as the Quilt Project provide forums for the exchange of information on network acceptance testing. These forums also may be useful for information sharing on Layer 1/Layer 2 analysis.

Needed operations and management capabilities include:

- Improved views of what is happening within and across the network
- Improved resource allocation and traffic management
- Capacity management
- Tools for fault isolation (looking for breaks, leaks, obstructions in the pipe)
- Measurement of end-to-end performance
- Tools for application performance monitoring
- Measurement for and as research
- Testing (10 Gbps and beyond, lambda testing, export and flow analysis)

2.9 Cost Models and Building Funding Support

ONT2 participants from operational research networks identified the sharing of information for business models and the need to build support for funding as important issues. To be effective, research funding needs to be sufficient and sustainable. However, in the current Federal climate, funds for research are limited so that users need to carry more of the financial burden and find ways to leverage their research expenditures.

Cost models are needed to identify scenarios for network financial sustainability. Such models would also help make the business case for sustained R&D investment in ONTs. The Quilt Project provides an existing model for demand aggregation, contract sharing, pricing information sharing, and service sharing (including IP bearer service, shared NOC services, and the swapping of light paths). IP transit sharing could help networks save on costs of providing commodity IP service and reduce the overhead and time spent on piecemeal peering.

Cost recovery is a concern of several networks, including GÉANT2, whose managers mentioned that while they experience marginal costs for light paths over an existing network, certain costs (e.g., transponders) are passed directly on to users and do not impact the financial situation of the network provider. This model encourages the use of light paths. For those providers that permit nonmembers to have their own light paths, the fees are often higher than those for members. This helps cover shared costs of the network not recovered by the charges for transponders and other equipment costs.

Research networks also need to coordinate outreach activities to inform members of the legislative and executive branches of government about the development, improvement, and increasing importance of networks and network research – with emphasis on how these networks provide engines for economic development and contribute to maintaining U.S. competitive advantages. With the growth of computational science, its reliance on high-performance networks, and its use by industry, continued improvements to the Internet are essential for economic development. Outreach activities might include conducting demonstrations and briefings, providing educational brochures, and the like. *Computational Science: Ensuring*

ONT2 Workshop Report

America's Competitiveness, the 2005 report by the President's Information Technology Advisory Committee (PITAC), may help in this area. To ensure that the latest innovations, tools, and techniques are presented, the research networks should coordinate with industry representatives.

Identified needs in research costing and funding support include:

- Costs, cost recovery, pricing (business models)
- Sharing of information to leverage resources
- Security costs
- Educational outreach/literature to impact resource decisions

2.9.1 Recommended Actions on Cost Models and Funding Support

- *Research networks should expand their cooperation and collaboration to enhance idea exchange, reduce costs, and increase the effectiveness of the limited research resources.*
- *Research networks should coordinate to provide outreach and to educate research funders.*

2.10 MorphNet Architecture to Support Transition From R&D to Production

Workshop participants identified the need to migrate experimental research technologies onto production networks but noted the lack of a well-defined process to do so. For example, DOE/SC has two networks, USN and ESnet. How should technologies be migrated from USN to ESnet without disrupting existing production traffic?

One approach, MorphNet architecture, provides support for multiple, concurrent, multi-layer views of the network (see: www.anl.gov/ECT/Public/research/morphnet.html). MorphNet architectures, such as NLR, can address potential conflicts in Federal research networks supporting both research and production activities over one network. GÉANT2 and some RONS are also built on the MorphNet model. For example, GÉANT2 provides Layer 1 optical wavelengths to researchers, while operating a production Layer 2/Layer 3 network on other wavelengths. An alternative solution is to provide Layer 2 Ethernet VLANs to researchers, while operational network business is operated on other VLANs at Layer 2 or 3. Participants said further research is needed to develop the capabilities of MorphNet architectures.

ONT2 Workshop Report

Appendix A: Workshop Administration and Attendees

Workshop Chairs

Joe Mambretti, Northwestern University
Dick desJardins, NASA/NREN

Panel Chairs

Panel A

Bill Wing, ORNL
Rick Summerhill, Internet2

Panel B

Guru Parulkar, NSF
Tom DeFanti, StarLight

Panel C

Steve Cotter, Internet2
Phil Dykstra, DREN: Ware on Earth

Panel D

Kevin Jones, NASA/NREN
Mark Johnson, MCNC

Panel E

Guru Parulkar, NSF

Panel F

Joe Mambretti, Northwestern University
Paul Love, Internet2

ONT2 Workshop Report

Attendees

Jules Aronson, National Library of Medicine (NLM)
William Banyai, Glimmerglass
Scott Beaudoin, WilTel Communications
Lou Berger, Movaz Networks, Inc
Joseph Berthold, CIENA
Dan Blumenthal, University of California at Santa Barbara
John Bowers, University of California at Santa Barbara
Maarten Büchli, DANTE
Larry Chao, NASA Ames Research Center (NREN)
James Cook, HPCMPO/DREN
Steve Cotter, Internet2
Cas D'Angelo, SoX/SLR, Georgia Institute of Technology
Hank Dardy, GIG-EF
Paul Daspit, Nortel/OMNinet
Tom DeFanti, OptIPuter, Translight, StarLight
Thomas DiPasquale, Nortel Networks
Phillip Dykstra, Ware On Earth/DREN
Dave Farber, Carnegie Mellon University, NLR
Ken Freeman, NASA Ames Research Center (NREN)
Denis Gallant, Meriton
Pat Gary, NASA Goddard Space Flight Center (GSFC)
Doug Gatchell, NSF
Chin Guok, ESnet
Ibrahim Habib, CHEETAH, City University of New York
Dave Hartzell, NASA Ames Research Center (NREN)
Koichi Hiragami, National Institute of Information & Communications Technology
Julio Ibarra, Florida International University
John Jamison, Cisco/NLR
Richard Jensen, U.S. Geological Survey (USGS)
Mark Johnson, MCNC
Kevin Jones, NASA Ames Research Center (NREN)
Wesley Kaplow, Qwest
Gigi Karmous-Edwards, MCNC
Kunihiro Kato, National Institute of Information & Communications Technology
Peter Kaufmann, DFN
Gordon Knoble, NASA GSFC
Joseph Lappa, Carnegie Mellon University, Pittsburgh Supercomputing Center
Carol Lawson, USGS
George Lee, Caltech
Tom Lehman, USC/ISI (DRAGON)
Tony Lisotta, Computer Sciences Corporation, NASA Stennis Space Center
Paul Love, Internet2
Joe Mambretti, International Center for Advanced Internet Research, Northwestern University
Jim McCabe, NASA/HQ/KAA
John McCoy, Pacific Northwest National Laboratory

ONT2 Workshop Report

John McDougale, NASA
Kevin McGrattan, Cisco/NLR
Joerg Micheel, Kaimai Networks
Grant Miller, NCO/NITRD
Biswanath Mukherjee, UC Davis
Pat Murphy, USGS
Yuko Nagaki, Northwestern University
Radha Nagarajan, Infinera Corporation (Panel F)
Kees Neggers, SURFnet/GLIF
Bill Owens, NYSERNet
Craig Partridge, BBN
Guru Parulkar, NSF
Drew Perkins, Infinera Corporation
Mike Rechtenbaugh, USGS
David Reese, CENIC, NLR
David Richardson, University of Washington
Anne Richeson, Qwest Government Services
James Sauer, Cisco Systems
George Seweryniak, DOE/SC
John Silvester, University of Southern California
Jerry Sobieski, Dragon/MAX; University of Maryland
Bill St. Arnaud, CANARIE/CA*net4
Jim Sugg, Juniper
Rick Summerhill, Internet2
Jerry Tounge, NASA Ames Research Center (NREN)
William Turnbull, NOAA
George Uhl, NASA GSFC
John Verduzco, L3
Tom West, NLR
Bessie Whitaker, NASA Marshall Space Flight Center (MSFC)
Ken White, NASA/Nasa Integrated Services Network (NISN)
Rodney Wilson, OmniNet, Nortel Networks
Bill Wing, DOE/SC UltraScience Net, Oak Ridge National Laboratory (ORNL)
Matt Zekauskas, Internet2

Appendix B: About LSN and the NITRD Program

Large Scale Networking (LSN) is one of the eight research areas – called program Component Areas (PCAs) – of the Federal government’s Networking and Information Technology Research and Development (NITRD) Program. The \$3-billion (President’s FY 2007 Budget request) NITRD activity is a collaborative enterprise of 14 Federal agencies that represents the Government’s main R&D investment portfolio in advanced computing, networking, software, and other information technologies. The NITRD member agencies are:

AHRQ – Agency for Healthcare Research and Quality
DARPA – Defense Advanced Research Projects Agency
DHS – Department of Homeland Security
DOE/NNSA – Department of Energy/National Nuclear Security Administration
DOE/SC – Department of Energy/Office of Science
EPA – Environmental Protection Agency
NARA – National Archives and Records Administration
NASA – National Aeronautics and Space Administration
NIH – National Institutes of Health
NIST – National Institute of Standards and Technology
NOAA – National Oceanic and Atmospheric Administration
NSF – National Science Foundation
NSA – National Security Agency
OSD and Service research organizations – Office of the Secretary of Defense and DoD Air Force, Army, and Navy research organizations

Representatives of the NITRD agencies that participate in LSN – NSF, OSD and DoD Service research organizations, NIH, DARPA, DOE/SC, NSA, NASA, AHRQ, NIST, DOE/NNSA, and NOAA – work together in the LSN Coordinating Group to coordinate Federal agency networking R&D in leading-edge networking technologies, services, and enhanced performance, including programs in new architectures, optical network testbeds, security, infrastructure, middleware, end-to-end performance measurement, and advanced network components; grid and collaboration networking tools and services; and engineering, management, and use of large-scale networks for scientific and applications R&D.

The Joint Engineering Team (JET) is one of three teams reporting to the LSN Coordinating Group:

- **The JET** coordinates the network architecture, connectivity, exchange points, and cooperation among Federal agency networks and other high-performance research networks, and provides close coordination of connectivity, interoperability, and services among government, academia, and industry to improve end-to-end user performance and avoid duplication of resources and efforts. The JET also coordinates international connectivity and interoperability.
- **The Middleware And Grid Infrastructure Coordination (MAGIC) Team** coordinates cooperation among Federal agencies, researchers, and commercial entities to research, develop, widely deploy, and use interoperable grid and middleware technologies, tools, and services and to provide a forum for international coordination.

ONT2 Workshop Report

- **The Networking Research Team (NRT)** coordinates agency networking research programs and shares networking research information among Federal agencies. It provides outreach to end users by disseminating networking research information and coordinating activities among applications developers and end users.

The other NITRD PCAs, each of which is also coordinated by an Interagency Working Group (IWG) or a Coordinating Group (CG) of agency representatives, are:

- **High-End Computing Infrastructure and Applications (HEC I&A)**
- **High-End Computing Research and Development (HEC R&D)**
- **Cyber Security and Information Assurance (CSIA)**
- **Human-Computer Interaction and Information Management (HCI&IM)**
- **High-Confidence Software and Systems (HCSS)**
- **Social, Economic, and Workforce Implications of IT and IT Workforce Development (SEW)**
- **Software Design and Productivity (SDP)**

In addition to the NITRD member agencies, a number of other Federal agencies also participate in the planning and coordination activities of NITRD's IWGs and CGs.

The NITRD Program is authorized by Congress under the High-Performance Computing (HPC) Act of 1991 (P.L. 102-194) and the Next Generation Internet Research Act of 1998 (P.L. 105-305). The strategic goals of the Program are to:

- Provide research and development foundations for assuring continued U.S. technological leadership in advanced networking, computing systems, software, and associated information technologies
- Provide research and development foundations for meeting the needs of the Federal government for advanced networking, computing systems, software, and associated information technologies
- Accelerate development and deployment of these technologies in order to maintain world leadership in science and engineering; enhance national defense and national and homeland security; improve U.S. productivity and competitiveness and promote long-term economic growth; improve the health of the U.S. citizenry; protect the environment; improve education, training, and lifelong learning; and improve the quality of life

Overall NITRD Program coordination is carried out by the Subcommittee on Networking and Information Technology Research and Development, under the aegis of the Committee on Technology of the National Science and Technology Council (NSTC). The Cabinet-level NSTC is the principal means by which the President coordinates the diverse science and technology programs across the Federal government. The National Coordination Office for the NITRD Program (NCO/NITRD) provides technical, planning, and budgetary support for the NITRD Subcommittee and activities in the NITRD PCAs. The NCO/NITRD also supports the Networking and Information Technology Subcommittee of the President's Council of Advisors on Science and Technology (PCAST).

ONT2 Workshop Report

Abbreviations/Acronyms

AAA – Authentication, authorization, and accounting

ATDnet – DARPA's Advanced Technology Demonstration Network

AMPATH – AMericaSPATH

API – Application program interface

ATM – Asynchronous transfer mode

BGP – Border Gateway Protocol

BNF – Backus-Naur Form

BOSSnet – Massachusetts Institute of Technology's BoSton-South optical network

BRUW – Internet2's Bandwidth Reservation for User Work project

CA*net4 – The network infrastructure of CANARIE

CANARIE – Canada's advanced Internet development organization

CENIC LA PoP – Corporation for Network Initiatives in California's Los Angeles Point of Presence

CER – Common event record

CERN – European Laboratory for Particle Physics

CISE – NSF's Computer and Information Science Engineering directorate

CISP – Common Internet Service Provisioning

CHEETAH – NSF's Circuit-switched High-speed End-to-End Transport Architecture project

COTS – Commercial off the shelf

DARPA – Defense Advanced Research Projects Agency

DoD – Department of Defense

DOE – Department of Energy

DOE/SC – Department of Energy Office of Science

DRAGON – NSF's Dynamic Resource Allocation via GMPLS Optical Networks project

DREN – DoD's Defense Research and Engineering Network

E-wave – An optical wavelength network link

EDDY – NSF's End-to-end Diagnostic DiscoverY project

ESnet – DOE's Energy Sciences network

GÉANT2 – Pan-European research and education network

Gbps – Gigabits per second

GENI – NSF's Global Environment for Networking Investigations

GGF – Global Grid Forum

GIG-EF – DoD's Global Information Grid Evaluation Facilities

GigaPoP – Gigabit per second Point of Presence

GLIF – Global Lambda Integrated Facility

GMPLS – Generalized Multi-Protocol Label Switching

GOLE – Global Open Lambda Exchange

HDXes – Nortel's high-density cross-connect optical switches

HOPI – Internet2's Hybrid Optical Packet Internet project

I/O – Input/output

IETF – Internet Engineering Task Force

iGRID – Annual international grid workshop

InfiniBand – A point-to-point high-speed switch fabric interconnect architecture

IP – Internet Protocol

IP/BGP – Internet Protocol/Border Gateway Protocol

ONT2 Workshop Report

IPv6 – Internet Protocol version 6	OMB – Office of Management and Budget
ISP – Internet service provider	ONT – Optical network testbed
IT – Information technology	OPN – Optical private network
ITU – International Telecommunications Union	OSCARS – DOE/SC’s On-demand Secure Circuits and Advanced Reservation System project
JET – LSN Coordinating Group's Joint Engineering Team	PCAST – President’s Council of Advisors on Science and Technology
KRLight – Global Open Lambda Exchange in Seoul, South Korea	PITAC – President's Information Technology Advisory Committee
LHC – Large Hadron Collider	PoP – Point of Presence
LSN – Large Scale Networking	QoS – Quality of service
LUSTRE – An open source distributed file system	R&D – Research and development
MAN LAN – Manhattan Landing exchange point	R&E – Research and education
MAX – Mid-Atlantic Exchange	RON – Regional optical network
Mbps – Megabits per second	SIP – Session Initiated Protocol
MPLS – Multi-Protocol Label Switching	SOA – Service-oriented architecture
NASA – National Aeronautics and Space Administration	SOX – South Atlantic Optical Exchange
NCO/NITRD – National Coordination Office for the NITRD Program	StarLight – NSF-funded international optical exchange point
NGIX – Next Generation Internet Exchange	SURFnet – Netherlands’ research and education network
NGN – Next-Generation Network	T-Lex – Tokyo Lambda Exchange
NITRD – The Federal Networking and Information Technology Research and Development Program	TCP/IP – Transmission Control Protocol/Internet Protocol
NLR – National LambdaRail	TeraGrid – NSF’s grid computing initiative
NOC – Network operations center	UCLP – User-controlled light path
NREN – NASA's Research and Education Network	UKLight – United Kingdom optical optical research network
NRENs – National research and education networks	USN – DOE/SC’s UltraScience Network
NSA – National Security Agency	VLAN – Virtual local area network
NSF – National Science Foundation	WAN – Wide area network
O&M – Operations and management	WG – Working Group
OC-192 – Optical Carrier rate of nearly 10 gigabits per second	WSRF – Web Services Resource Framework
OEP – Open Exchange Point	XML – eXtensible Markup Language
OIF – Optical Internetworking Forum	

Suite II-405
4201 Wilson Blvd.
Arlington, VA 22230

Phone : (703) 292-4873 Fax : (703) 292-9097

Web address: <http://www.nitrd.gov>

E-mail: nco@nitrd.gov

National Coordination Office for
Networking and Information Technology
Research and Development

